

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L3	16	709/203 and ring and load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:51
L4	2	709/203 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:53
L5	1	709/217 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:54
L6	0	709/218 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:54
L7	0	709/219 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:54
L8	2	709/226 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:55
L9	1	709/235 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:55
L10	2	709/238 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:55
L11	0	709/239 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:55
L12	1	709/245 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:56

L13	1	718/105 and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:56
L14	7	ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:58
L15	1	"718"/\$ and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:57
L16	4	"709"/\$ and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:57
L17	3	"710"/\$ and ring same load near3 balanc\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:57
L18	1	ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413") and @pd>"20040701"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 08:59
L19	14	ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 10:33
L20	3	"714"/\$ and ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 10:37
L21	3	714/1-5.ccls. and ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 10:38
L22	0	714/10-11.ccls. and ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 10:38

L23	0	710/3-4.ccls. and ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 10:38
L24	1	709/245 and ring same load near3 balanc\$4 same fail\$4 and (@ad<"20010413" @rlad<"20010413")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/16 10:39


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[ring](#) and [fail](#) and [server](#) and [address](#) and [name](#) and [translat](#)

Found 50,559 of 145,831

 Sort results by
[Save results to a Binder](#)

 Display results
[Search Tips](#)
 Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [File servers for network-based distributed systems](#)

Liba Svobodova

 December 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 4

 Full text available: [pdf\(4.23 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)


2 [Address translation in telecommunication features](#)

Pamela Zave

 January 2004 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 13 Issue 1

 Full text available: [pdf\(378.36 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Address translation causes a wide variety of interactions among telecommunication features. This article begins with a formal model of address translation and its effects, and with principles for understanding how features should interact in the presence of address translation. There is a simple and intuitive set of constraints on feature behavior so that features will interact according to the principles. This scheme (called "ideal address translation") has provable properties, is modular (expl ...)

Keywords: Component architecture, feature interaction, formal methods, network addressing, network protocols, network security, requirements, telecommunications

3 [Distributed operating systems](#)

Andrew S. Tanenbaum, Robbert Van Renesse

 December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

 Full text available: [pdf\(5.49 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

4 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4Full text available:  pdf(5.33 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

5 An open architecture for next-generation telecommunication services

Gregory W. Bond, Eric Cheung, K. Hal Purdy, Pamela Zave, J. Christopher Ramming

February 2004 **ACM Transactions on Internet Technology (TOIT)**, Volume 4 Issue 1Full text available:  pdf(237.24 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An open (in the sense of extensible and programmable) architecture for IP telecommunications must be based on a comprehensive strategy for managing feature interaction. We describe our experience with BoxOS, an IP telecommunication platform that implements the DFC technology for feature composition. We present solutions to problems, common to all efforts in IP telecommunications, of feature distribution, interoperability, and media management. We also explain how BoxOS addresses many deficiencies ...

Keywords: Component architectures, Intelligent Network architecture, Session Initiation Protocol, electronic mail, feature interaction, instant messaging, multimedia systems, network addressing, network interoperation, network optimization, network protocols, service creation

6 The file system of an integrated local network

Paul J. Leach, Paul H. Levine, James A. Hamilton, Bernard L. Stumpf

March 1985 **Proceedings of the 1985 ACM thirteenth annual conference on Computer Science**Full text available:  pdf(1.78 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The distributed file system component of the DOMAIN system is described. The DOMAIN system is an architecture for networks of personal workstations and servers which creates an integrated distributed computing environment. The distinctive features of the file system include: objects addressed by unique identifiers (UIDs); transparent access to objects, regardless of their location in the network; the abstraction of a single level store for accessing all objects; and the layering of a network ...

7 Peer to peer networks: Tarzan: a peer-to-peer anonymizing network layer

Michael J. Freedman, Robert Morris

November 2002 **Proceedings of the 9th ACM conference on Computer and communications security**Full text available:  pdf(242.72 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Tarzan is a peer-to-peer anonymous IP network overlay. Because it provides IP service, Tarzan is general-purpose and transparent to applications. Organized as a decentralized peer-to-peer overlay, Tarzan is fault-tolerant, highly scalable, and easy to manage. Tarzan achieves its anonymity with layered encryption and multi-hop routing, much like a Chaumian mix. A message initiator chooses a path of peers pseudo-randomly through a

restricted topology in a way that adversaries cannot easily influenc ...

Keywords: IP tunnels, anonymity, cover traffic, distributed trust, mix-nets, overlay networks, peer-to-peer

8 A comparison of two network-based file servers

James G. Mitchell, Jeremy Dion

April 1982 **Communications of the ACM**, Volume 25 Issue 4

Full text available:  pdf(1.50 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper compares two working network-based file servers, the Xerox Distributed File System (XDFS) implemented at the Xerox Palo Alto Research Center, and the Cambridge File Server (CFS) implemented at the Cambridge University Computer Laboratory. Both servers support concurrent random access to files using atomic transactions, both are connected to local area networks, and both have been in service long enough to enable us to draw lessons from them for future file servers. We ...

9 Session summaries from the 17th symposium on operating systems principle (SOSP'99)

Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Full text available:  pdf(3.15 MB)

Additional Information: [full citation](#), [index terms](#)

10 On randomization in sequential and distributed algorithms

Rajiv Gupta, Scott A. Smolka, Shaji Bhaskar

March 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 1

Full text available:  pdf(8.01 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Probabilistic, or randomized, algorithms are fast becoming as commonplace as conventional deterministic algorithms. This survey presents five techniques that have been widely used in the design of randomized algorithms. These techniques are illustrated using 12 randomized algorithms—both sequential and distributed—that span a wide range of applications, including: primality testing (a classical problem in number theory), interactive probabilistic proof s ...

Keywords: Byzantine agreement, CSP, analysis of algorithms, computational complexity, dining philosophers problem, distributed algorithms, graph isomorphism, hashing, interactive probabilistic proof systems, leader election, message routing, nearest-neighbors problem, perfect hashing, primality testing, probabilistic techniques, randomized or probabilistic algorithms, randomized quicksort, sequential algorithms, transitive tournaments, universal hashing

11 Generative communication in Linda

David Gelernter

January 1985 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 7 Issue 1

Full text available:  pdf(2.48 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Generative communication is the basis of a new distributed programming language that is intended for systems programming in distributed settings generally and on integrated

network computers in particular. It differs from previous interprocess communication models in specifying that messages be added in tuple-structured form to the computation environment, where they exist as named, independent entities until some process chooses to receive them. Generative communication results in a number ...

12 Integrating security in a large distributed system

M. Satyanarayanan

August 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 3

Full text available:  pdf(2.90 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Andrew is a distributed computing environment that is a synthesis of the personal computing and timesharing paradigms. When mature, it is expected to encompass over 5,000 workstations spanning the Carnegie Mellon University campus. This paper examines the security issues that arise in such an environment and describes the mechanisms that have been developed to address them. These mechanisms include the logical and physical separation of servers and clients, support for secure communication ...

13 Programming languages for distributed computing systems

Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum

September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Full text available:  pdf(6.50 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

14 Grapevine: an exercise in distributed computing

Andrew D. Birrell, Roy Levin, Michael D. Schroeder, Roger M. Needham

April 1982 **Communications of the ACM**, Volume 25 Issue 4

Full text available:  pdf(1.71 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Grapevine is a multicomputer system on the Xerox research internet. It provides facilities for the delivery of digital messages such as computer mail; for naming people, machines, and services; for authenticating people and machines; and for locating services on the internet. This paper has two goals: to describe the system itself and to serve as a case study of a real application of distributed computing. Part I describes the set of services provided by Grapevine and how its data and funct ...

15 Special issue on persistent object systems: Orthogonally persistent object systems

Malcolm Atkinson, Ronald Morrison

July 1995 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 4 Issue 3

Full text available:  pdf(5.02 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Persistent Application Systems (PASs) are of increasing social and economic importance. They have the potential to be long-lived, concurrently accessed, and consist of large bodies of data and programs. Typical examples of PASs are CAD/CAM systems, office automation, CASE tools, software engineering environments, and patient-care support systems in hospitals. Orthogonally persistent object systems are intended to provide improved support for the design, construction, maintenance, and operation o ...

Keywords: database programming languages, orthogonal persistence, persistent application systems, persistent programming languages

16 The family of concurrent logic programming languages

Ehud Shapiro

September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Full text available:  pdf(9.62 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Concurrent logic languages are high-level programming languages for parallel and distributed systems that offer a wide range of both known and novel concurrent programming techniques. Being logic programming languages, they preserve many advantages of the abstract logic programming model, including the logical reading of programs and computations, the convenience of representing data structures with logical terms and manipulating them using unification, and the amenability to metaprogrammin ...

17 A scalable content-addressable network

Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, Scott Schenker

August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications**, Volume 31 Issue 4

Full text available:  pdf(155.64 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

18 Virtual machine monitors: Xen and the art of virtualization

Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, Andrew Warfield

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Full text available:  pdf(168.76 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Numerous systems have been designed which use virtualization to subdivide the ample resources of a modern computer. Some require specialized hardware, or cannot support commodity operating systems. Some target 100% binary compatibility at the expense of performance. Others sacrifice security or functionality for speed. Few offer resource isolation or performance guarantees; most provide only best-effort provisioning, risking denial of service. This paper presents Xen, an x86 virtual machine monit ...

Keywords: hypervisors, paravirtualization, virtual machine monitors

19 Local networking and internetworking in the V-system

David R. Cheriton

October 1983 **ACM SIGCOMM Computer Communication Review , Proceedings of the eighth symposium on Data communications**, Volume 13 Issue 4

Full text available:  pdf(896.12 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Local networking can be treated as a subset of internetworking for remote terminal access and file transfer. However, a distributed operating system, such as the V-System uses a local network more as an extended backplane than a fast, miniature long-haul network. This paper describes the use of server-based "intelligent gateways" to provide internetworking using standard protocols in conjunction with an efficient light-weight protocol for V IPC on a lo ...

**20** Client-server computing

Alok Sinha

July 1992 **Communications of the ACM**, Volume 35 Issue 7

Full text available: pdf(7.53 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#),
[review](#)**Keywords:** client-server computing

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: The ACM Digital Library The Guide

THE ACM DIGITAL LIBRARY
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used
[ring](#) and [fail](#) and [server](#) and [address](#) and [name](#) and [translat](#) and [load balanc](#)

Found 52,005 of 145,831

 Sort results by
[Save results to a Binder](#)
[Try an Advanced Search](#)

 Display results
[Search Tips](#)
[Try this search in The ACM Guide](#)
 Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 Distributed operating systems

Andrew S. Tanenbaum, Robbert Van Renesse

 December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

 Full text available: [pdf\(5.49 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

2 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

 December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4

 Full text available: [pdf\(5.33 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

3 The family of concurrent logic programming languages

Ehud Shapiro

 September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

 Full text available: [pdf\(9.62 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Concurrent logic languages are high-level programming languages for parallel and distributed systems that offer a wide range of both known and novel concurrent programming techniques. Being logic programming languages, they preserve many advantages of the abstract logic programming model, including the logical reading of programs and computations, the convenience of representing data structures with logical

terms and manipulating them using unification, and the amenability to metaprogrammin ...

4 Load balancing and fault tolerance in workstation clusters migrating groups of communicating processes

S. Petri, H. Langendorfer

October 1995 **ACM SIGOPS Operating Systems Review**, Volume 29 Issue 4

Full text available:  pdf(894.43 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

In the past, several process migration facilities for distributed systems have been developed. Due to the complex nature of the subject, all those facilities have limitations that make them usable for only limited classes of applications and environments. We discuss some of the usual limitations and possible solutions. Specifically, we focus on migration of groups of collaborating processes between Unix systems without kernel modifications, and from this we derive the design for a migration syst ...

5 A scalable content-addressable network

Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, Scott Schenker

August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications**, Volume 31 Issue 4

Full text available:  pdf(155.64 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Session summaries from the 17th symposium on operating systems principle (SOSP'99)

Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Full text available:  pdf(3.15 MB) Additional Information: [full citation](#), [index terms](#)

7 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available:  pdf(9.37 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

8 Special issue on persistent object systems: Orthogonally persistent object systems

Malcolm Atkinson, Ronald Morrison

July 1995 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 4 Issue 3

Full text available: Additional Information:

 pdf(5.02 MB)[full citation, abstract, references, citations](#)

Persistent Application Systems (PASs) are of increasing social and economic importance. They have the potential to be long-lived, concurrently accessed, and consist of large bodies of data and programs. Typical examples of PASs are CAD/CAM systems, office automation, CASE tools, software engineering environments, and patient-care support systems in hospitals. Orthogonally persistent object systems are intended to provide improved support for the design, construction, maintenance, and operation o ...

Keywords: database programming languages, orthogonal persistence, persistent application systems, persistent programming languages

9 An open architecture for next-generation telecommunication services 

Gregory W. Bond, Eric Cheung, K. Hal Purdy, Pamela Zave, J. Christopher Ramming
February 2004 **ACM Transactions on Internet Technology (TOIT)**, Volume 4 Issue 1

Full text available:  pdf(237.24 KB) Additional Information: [full citation, abstract, references, index terms](#)

An open (in the sense of extensible and programmable) architecture for IP telecommunications must be based on a comprehensive strategy for managing feature interaction. We describe our experience with BoxOS, an IP telecommunication platform that implements the DFC technology for feature composition. We present solutions to problems, common to all efforts in IP telecommunications, of feature distribution, interoperability, and media management. We also explain how BoxOS addresses many deficiencies ...

Keywords: Component architectures, Intelligent Network architecture, Session Initiation Protocol, electronic mail, feature interaction, instant messaging, multimedia systems, network addressing, network interoperation, network optimization, network protocols, service creation

10 Peer to peer networks: Tarzan: a peer-to-peer anonymizing network layer 

Michael J. Freedman, Robert Morris

November 2002 **Proceedings of the 9th ACM conference on Computer and communications security**

Full text available:  pdf(242.72 KB) Additional Information: [full citation, abstract, references, citations, index terms](#)

Tarzan is a peer-to-peer anonymous IP network overlay. Because it provides IP service, Tarzan is general-purpose and transparent to applications. Organized as a decentralized peer-to-peer overlay, Tarzan is fault-tolerant, highly scalable, and easy to manage. Tarzan achieves its anonymity with layered encryption and multi-hop routing, much like a Chaumian mix. A message initiator chooses a path of peers pseudo-randomly through a restricted topology in a way that adversaries cannot easily influence ...

Keywords: IP tunnels, anonymity, cover traffic, distributed trust, mix-nets, overlay networks, peer-to-peer

11 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors 

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available:  pdf(1.93 MB) Additional Information: [full citation, abstract, references, citations, index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a syste ...

12 Scale and performance in a distributed file system

John H. Howard, Michael L. Kazar, Sherri G. Menees, David A. Nichols, M. Satyanarayanan, Robert N. Sidebotham, Michael J. West

February 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 1

Full text available:  pdf(2.38 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The Andrew File System is a location-transparent distributed file system that will eventually span more than 5000 workstations at Carnegie Mellon University. Large scale affects performance and complicates system operation. In this paper we present observations of a prototype implementation, motivate changes in the areas of cache validation, server process structure, name translation, and low-level storage representation, and quantitatively demonstrate Andrews ability to scale gracefully. W ...

13 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson

June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Full text available:  pdf(3.60 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

14 On randomization in sequential and distributed algorithms

Rajiv Gupta, Scott A. Smolka, Shaji Bhaskar

March 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 1

Full text available:  pdf(8.01 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Probabilistic, or randomized, algorithms are fast becoming as commonplace as conventional deterministic algorithms. This survey presents five techniques that have been widely used in the design of randomized algorithms. These techniques are illustrated using 12 randomized algorithms—both sequential and distributed—that span a wide range of applications, including primality testing (a classical problem in number theory), interactive probabilistic proof s ...

Keywords: Byzantine agreement, CSP, analysis of algorithms, computational complexity, dining philosophers problem, distributed algorithms, graph isomorphism, hashing, interactive probabilistic proof systems, leader election, message routing, nearest-neighbors problem, perfect hashing, primality testing, probabilistic techniques, randomized or probabilistic algorithms, randomized quicksort, sequential algorithms, transitive tournaments, universal hashing

15 Virtual machine monitors: Xen and the art of virtualization

Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, Andrew Warfield

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Full text available:  [pdf\(168.76 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Numerous systems have been designed which use virtualization to subdivide the ample resources of a modern computer. Some require specialized hardware, or cannot support commodity operating systems. Some target 100% binary compatibility at the expense of performance. Others sacrifice security or functionality for speed. Few offer resource isolation or performance guarantees; most provide only best-effort provisioning, risking denial of service. This paper presents Xen, an x86 virtual machine monitor ...

Keywords: hypervisors, paravirtualization, virtual machine monitors

16 Node autonomy in distributed systems

H. Garcia Molina, B. Kogan

January 2000 **Proceedings of the first international symposium on Databases in parallel and distributed systems**

Full text available:  [pdf\(925.19 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The goal of this paper is to explore the notion of node autonomy in distributed computer systems. Some motivations for autonomy are exposed. Different facets of autonomy as well as relationships among them are discussed. Finally, we look into how autonomy affects other aspects of distributed computing, including timeliness, correctness, load sharing, data sharing, and data replication.

17 Programming languages for distributed computing systems

Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum

September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Full text available:  [pdf\(6.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

18 Fast and flexible application-level networking on exokernel systems

Gregory R. Ganger, Dawson R. Engler, M. Frans Kaashoek, Héctor M. Briceño, Russell Hunt, Thomas Pinckney

February 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 1

Full text available:  [pdf\(500.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Application-level networking is a promising software organization for improving performance and functionality for important network services. The Xok/ExOS exokernel system includes application-level support for standard network services, while at the same time allowing application writers to specialize networking services. This paper describes how Xok/ExOS's kernel mechanisms and library operating system organization achieve this flexibility, and

retrospectively shares our experiences an ...

Keywords: Extensible systems, OS structure, fast servers, network services

19 Generative communication in Linda

David Gelernter

January 1985 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,
Volume 7 Issue 1

Full text available:  pdf(2.48 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Generative communication is the basis of a new distributed programming language that is intended for systems programming in distributed settings generally and on integrated network computers in particular. It differs from previous interprocess communication models in specifying that messages be added in tuple-structured form to the computation environment, where they exist as named, independent entities until some process chooses to receive them. Generative communication results in a number ...

20 Parallel execution of prolog programs: a survey

Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo

July 2001 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,
Volume 23 Issue 4

Full text available:  pdf(1.95 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Since the early days of logic programming, researchers in the field realized the potential for exploitation of parallelism present in the execution of logic programs. Their high-level nature, the presence of nondeterminism, and their referential transparency, among other characteristics, make logic programs interesting candidates for obtaining speedups through parallel execution. At the same time, the fact that the typical applications of logic programming frequently involve irregular computatio ...

Keywords: Automatic parallelization, constraint programming, logic programming, parallelism, prolog

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

load balanc* and fail* and neighbor* and ring

SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used load balanc and fail and neighbor and ring

Found 29,965 of 145,831

Sort results by relevance Save results to a BinderDisplay results expanded form Search Tips Open results in a new windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [A novel load balancing scheme for the tele-traffic hot spot problem in cellular networks](#)

Sajal K. Das, Sanjoy K. Sen, Rajeev Jayaram

July 1998 **Wireless Networks**, Volume 4 Issue 4Full text available: [pdf\(487.85 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We propose a dynamic load balancing scheme for the tele-traffic hot spot problem in cellular networks. A tele-traffic hot spot is a region of adjacent hot cells where the channel demand has exceeded a certain threshold. A hot spot is depicted as a stack of hexagonal 'Rings' of cells and is classified as complete if all cells within it are hot. Otherwise it is termed incomplete. The rings containing all cold cells outside the hot spot are called 'Peripheral Rings'. Our load balancing scheme ...

2 [On balancing the load in a clustered web farm](#)

Joel L. Wolf, Philip S. Yu

November 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 2Full text available: [pdf\(612.40 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this article we propose a novel, yet practical, scheme which attempts to optimally balance the load on the servers of a clustered Web farm. The goal in solving this performance problem is to achieve minimal average response time for customer requests, and thus ultimately achieve maximal customer throughput. The article decouples the overall problem into two related but distinct mathematical subproblems, one static and one dynamic. We believe this natural decoupling is one of the major contrib ...

Keywords: Clustered Web farms, combinatorial optimization, load balancing, resource allocation problems

3 [Peer-To-peer systems: Simple efficient load balancing algorithms for peer-to-peer systems](#)

David R. Karger, Matthias Ruhl

June 2004 **Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures**Full text available: [pdf\(135.23 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Load balancing is a critical issue for the efficient operation of peer-to-peer networks. We

give two new load-balancing protocols whose provable performance guarantees are within a constant factor of optimal. Our protocols refine the *consistent hashing* data structure that underlies the Chord (and Koorde) P2P network. Both preserve Chord's logarithmic query time and near-optimal data migration cost. Consistent hashing is an instance of the distributed hash table (DHT) paradigm for assigning ...

Keywords: load balancing, peer-to-peer systems

4 Parallel interval-Newton using message passing: dynamic load balancing strategies

Chao-Yang Gau, Mark A. Statherr

November 2001 **Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:  pdf(296.07 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Branch-and-prune and branch-and-bound techniques are commonly used for intelligent search in finding all solutions, or the optimal solution, within a space of interest. The corresponding binary tree structure provides a natural parallelism allowing concurrent evaluation of subproblems using parallel computing technology. Of special interest here are techniques derived from interval analysis, in particular an interval-Newton/generalized-bisection procedure. In this context, we discuss issues of I ...

Keywords: branch-and-bound, branch-and-prune, global optimization, interval analysis, nonlinear equations, parallel computing

5 Internet applications: Load balancing and locality in range-queriable data structures

James Aspnes, Jonathan Kirsch, Arvind Krishnamurthy

July 2004 **Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing**

Full text available:  pdf(236.00 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We describe a load-balancing mechanism for assigning elements to servers in a distributed data structure that supports range queries. The mechanism ensures both load-balancing with respect to an arbitrary load measure specified by the user and geographical locality, assigning elements with similar keys to the same server. Though our mechanism is specifically designed to improve the performance of skip graphs, it can be adapted to provide deterministic, locality-preserving load-balancing to any d ...

Keywords: overlay networks, peer-to-peer systems

6 An analysis of diffusive load-balancing

Raghu Subramanian, Isaac D. Scherson

August 1994 **Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures**

Full text available:  pdf(575.66 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Diffusion is a well-known algorithm for load-balancing in which tasks move from heavily-loaded processors to lightly-loaded neighbors. This paper presents a rigorous analysis of the performance of the diffusion algorithm on arbitrary networks. It is shown that the running time of the diffusion algorithm is bounded by: $\Theta(\log N \cdot \Delta)$ \leq Time $\leq O(N \cdot \Delta^2)$, where N is the number of nodes ...

7 Job scheduling in rings

Perry Fizzano, David Karger, Clifford Stein, Joel Wein
August 1994 **Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures**

Full text available: [pdf\(1.15 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

8 Storage: GEM: graph eMbedding for routing and data-centric storage in sensor networks without geographic information

James Newsome, Dawn Song
November 2003 **Proceedings of the first international conference on Embedded networked sensor systems**

Full text available: [pdf\(273.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The widespread deployment of sensor networks is on the horizon. One of the main challenges in sensor networks is to process and aggregate data in the network rather than wasting energy by sending large amounts of raw data to reply to a query. Some efficient data dissemination methods, particularly data-centric storage and information aggregation, rely on efficient routing from one node to another. In this paper we introduce GEM (Graph EMbedding for sensor networks), an infrastructure for node-to ...

Keywords: data-centric storage, routing, sensor networks

9 On zone-balancing of peer-to-peer networks: analysis of random node join

Xiaoming Wang, Yueping Zhang, Xiafeng Li, Dmitri Loguinov
June 2004 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the joint international conference on Measurement and modeling of computer systems**, Volume 32 Issue 1

Full text available: [pdf\(282.76 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Balancing peer-to-peer graphs, including zone-size distributions, has recently become an important topic of peer-to-peer (P2P) research [1], [2], [6], [19], [31], [36]. To bring analytical understanding into the various peer-join mechanisms, we study how zone-balancing decisions made during the initial sampling of the peer space affect the resulting zone sizes and derive several asymptotic results for the maximum and minimum zone sizes that hold with high probability.

Keywords: balls-into-bins, load-balancing, modeling, peer-to-peer

10 Dynamic load balancing strategies for conservative parallel simulations

Azzedine Boukerche, Sajal K. Das
June 1997 **ACM SIGSIM Simulation Digest , Proceedings of the eleventh workshop on Parallel and distributed simulation**, Volume 27 Issue 1

Full text available: [pdf\(1.09 MB\)](#) [Publisher Site](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper studies the problem of load balancing for conservative parallel simulations for execution on a multicomputer. The synchronization protocol makes use of Chandy-Misra null-messages. We propose a dynamic load balancing algorithm which assumes no compile time knowledge about the workload parameters. It is based upon a process migration mechanism, and the notion of CPU-queue length, which indicates the workload at each processor. We examine two variations for the algorithm which we refer to ...

11 First and second order diffusive methods for rapid, coarse, distributed load balancing

(extended abstract)

Bhaskar Ghosh, S. Muthukrishnan, Martin H. Schultz

June 1996 **Proceedings of the eighth annual ACM symposium on Parallel algorithms and architectures**

Full text available: [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

12 Distributed operating systems

Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Full text available: [pdf\(5.49 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

13 Session 5C: Skip graphs

James Aspnes, Gauri Shah

January 2003 **Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms**

Full text available: [pdf\(974.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Skip graphs are a novel distributed data structure, based on skip lists, that provide the full functionality of a balanced tree in a distributed system where elements are stored in separate nodes that may fail at any time. They are designed for use in searching peer-to-peer networks, and by providing the ability to perform queries based on key ordering, they improve on existing search tools that provide only hash table functionality. Unlike skip lists or other tree data structures, skip graphs a ...

14 Distributed information systems: Mercury: supporting scalable multi-attribute range queries

Ashwin R. Bharambe, Mukesh Agrawal, Srinivasan Seshan

August 2004 **Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications**

Full text available: [pdf\(1.29 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents the design of Mercury, a scalable protocol for supporting multi-attribute range-based searches. Mercury differs from previous range-based query systems in that it supports *multiple attributes* as well as performs *explicit load balancing*. To guarantee efficient routing and load balancing, Mercury uses novel light-weight sampling mechanisms for uniformly sampling random nodes in a highly dynamic overlay network. Our evaluation shows that Mercury is able to achieve ...

Keywords: distributed hash tables, load balancing, peer-to-peer systems, random sampling, range queries

15 Multi-scale self-simulation: a technique for reconfiguring arrays with faults

Richard Cole, Bruce Maggs, Ramesh Sitaraman

June 1993 **Proceedings of the twenty-fifth annual ACM symposium on Theory of computing**

Full text available:  pdf(1.43 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

16 The multicast policy and its relationship to replicated data placement

Ouri Wolfson, Amir Milo

March 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 1

Full text available:  pdf(1.61 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In this paper we consider the communication complexity of maintaining the replicas of a logical data-item, in a database distributed over a computer network. We propose a new method, called the minimum spanning tree write, by which a processor in the network should multicast a write of a logical data-item, to all the processors that store replicas of the items. Then we show that the minimum spanning tree write is optimal from the communication cost point of view. We also demonstrate that the ...

Keywords: NP-Complete, complexity, computer network, file allocation, message passing

17 Routing I: Novel architectures for P2P applications: the continuous-discrete approach

Moni Naor, Udi Wieder

June 2003 **Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures**

Full text available:  pdf(260.13 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We propose a new approach for constructing P2P networks based on a dynamic decomposition of a continuous space into cells corresponding to processors. We demonstrate the power of these design rules by suggesting two new architectures, one for DHT (Distributed Hash Table) and the other for dynamic expander networks. The DHT network, which we call Distance Halving allows logarithmic routing and load, while preserving constant degrees. It offers an optimal tradeoff between the degree and the dilatation ...

Keywords: distributed systems, fault tolerance, hash tables, peer-to-peer

18 On randomization in sequential and distributed algorithms

Rajiv Gupta, Scott A. Smolka, Shaji Bhaskar

March 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 1

Full text available:  pdf(8.01 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Probabilistic, or randomized, algorithms are fast becoming as commonplace as conventional deterministic algorithms. This survey presents five techniques that have been widely used in the design of randomized algorithms. These techniques are illustrated using 12 randomized algorithms—both sequential and distributed—that span a wide range of applications, including: primality testing (a classical problem in number theory), interactive probabilistic proofs ...

Keywords: Byzantine agreement, CSP, analysis of algorithms, computational complexity, dining philosophers problem, distributed algorithms, graph isomorphism, hashing, interactive probabilistic proof systems, leader election, message routing, nearest-neighbors problem, perfect hashing, primality testing, probabilistic techniques, randomized or probabilistic algorithms, randomized quicksort, sequential algorithms, transitive tournaments, universal hashing

19 Chord: A scalable peer-to-peer lookup service for internet applications

Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, Hari Balakrishnan

August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications**, Volume 31 Issue 4Full text available:  pdf(205.73 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**20 A scalable content-addressable network**

Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, Scott Schenker

August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications**, Volume 31 Issue 4Full text available:  pdf(155.64 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[IEEE HOME](#) | [SEARCH IEEE](#) | [SHOP](#) | [WEB ACCOUNT](#) | [CONTACT IEEE](#)[Membership](#) [Publications/Services](#) [Standards](#) [Conferences](#) [Careers/Jobs](#)Welcome
United States Patent and Trademark Office[Help](#) [FAQ](#) [Terms](#) [IEEE Peer Review](#)[Quick Links](#)**Welcome to IEEE Xplore®**

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced
- CrossRef

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved

[IEEE HOME](#) | [SEARCH IEEE](#) | [SHOP](#) | [WEB ACCOUNT](#) | [CONTACT IEEE](#)[Membership](#) [Publications/Services](#) [Standards](#) [Conferences](#) [Careers/Jobs](#)**IEEE Xplore®**
RELEASE 1.8Welcome
United States Patent and Trademark Office[Help](#) [FAQ](#) [Terms](#) [IEEE Peer Review](#)**Quick Links****Welcome to IEEE Xplore®**

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced
- CrossRef

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

Print Format

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER: _____**

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.